ADVANCED RESEARCH PROJECTS AGENCY
Washington, D. C. 20301

Program Plan No. 723

Date: 3 June 1968

AO #1260

Program Title: RESOURCE SHARING COMPUTER NETWORKS

Type of Contractor: To be selected from 4 - Industry (Profit)

Project from which Funded: Information Processing Techniques (8D30)

Prepared by: Lawrence G. Roberts
Lawrence G. Roberts
Program Manager

Concurred in: Robert W. Taylor
Robert W. Taylor, Director for
Information Processing Techniques

Approved by: [signature]
Director, ARPA

Date:

Type of Work: B (Applied Research)

Field of Work: 77 (Computer Sciences)

ADVA. ɔED RESEARCH PROJECTS AL ɔNCY
Washington, D.C. 20301

Program Plan No. 723

Date: 3 June 1968

RESOURCE SHARING COMPUTER NETWORKS

A. Objective of the Program.

The objective of this program is twofold: (1) To develop techniques and obtain experience on interconnecting computers in such a way that a very broad class of interactions are possible, and (2) To improve and increase computer research productivity through resource sharing. By establishing a network tying IPT's research centers together, both goals are achieved. In fact, the most efficient way to develop the techniques needed for an effective network is by involving the research talent at these centers in prototype activity.

Just as time-shared computer systems have permitted groups of hundreds of individual users to share hardware and software resources with one another, networks connecting dozens of such systems will permit resource sharing between thousands of users. Each system, by virtue of being time-shared, can offer any of its services to another computer system on demand. The most important criterion for the type of network interconnection desired is that any user or program on any of the networked computers can utilize any program or subsystem available on any other computer without having to modify the remote program.

B. Technical Need and Background of the Program.

1. Scientific Environment.

Currently there are thousands of computer centers in the country, each of which operates almost completely autonomously. There is some trading of programs between those machines which are sufficiently similar to allow this, and there is technical communication through publications of technical meetings describing techniques developed. However, since the computer field is growing at such a rapid rate, a more immediate mechanism must be developed if there is to be significant cross-fertilization in sharing between these many centers. Although the same problem exists in many technological areas, the solution is most easily found and implemented by the computer community. If a sufficiently reliable and

capable network were established linking these centers, many improvements could be obtained. There would be less duplication of large programs and systems, some of which require hundreds of man months of effort. Currently such programs must be reprogrammed for each machine where they are needed even if they are only required occasionally. It is estimated that such duplicative efforts more than double the national costs of creating and maintaining the software. A network will not eliminate all of this duplication but can be used for those functions which are only infrequently called and those which only need to be tested. Further, there are large data files available at individual locations which are not valuable enough to warrant duplication at every computer center, but from which segments could be obtained at any network location. For example, within the ARPA research centers there are files of speech samples, digitized pictures and the semantic definitions of most English words.

Often it is important at a research establishment to test a new language developed at another installation to determine what features should be incorporated into local languages. Currently one either reprograms the language on his local machine or obtains sufficient remote console time to evaluate the language. Although it may be preferable to use the original system via remote consoles, this is often difficult or impossible due to console incompatibility. With an interactive network it is possible to use one's local consoles through the local computer to access the remote system, thus eliminating the need for compatible consoles and at the same time reducing the communications costs by several orders of magnitude.

Another important application of a network is to link specialized computers to general purpose computer centers. ILLIAC III and ILLIAC IV are examples of specialized machines designed to process special classes of problems. With recent improvements in the hardware area, it will become more cost effective to design and construct computers particularly efficient at specialized tasks (e.g., compiling, list processing and information retrieval). Making such machines available to all the computer research establishments would significantly increase the capability at these other centers.

2. <u>Military Environment</u>.

The military environment, like the scientific environment, includes thousands of computers of various vintages and vendors. The traditional staff elements (Personnel, Intelligence, Operations (Command and Control), Logistics and Communications) throughout all the Services

are using various machines with varying degrees of success. It is not unusual to find a greater continuity of community interest within, say, logistic commands of the Army, Navy and Air Force than that which exists across, say, one Air Force staff. With the current fractionation of computer resources in the absence of any technology permitting the interconnection and sharing of these resources, the current situation can only get worse. Military personnel trained to use one manufacturer's equipment must often be trained again to use a different manufacturer's equipment as the personnel move from one military station to another. Machines procured from different manufacturers require as many different user training programs as there are machines thus inhibiting positive transfer of training that could accumulate through the rotation of military personnel. Those data files and programs which have common utility to many military organizations and installations must be stored, created and maintained separately at each different machine. Military systems interconnected in a distributed interactive network obviate such constraints.

3. Previous Work.

Relatively little work has been done in the past on interactive computer networks and it is mainly with the advent of widespread time-sharing that such nets become feasible. Most previous work has concentrated on either load sharing or message handling goals. In 1963 ARPA initiated a project at UCLA to develop a network linking computers in the Western Data Processing Center. This work was aimed at enabling several similar computers to load-share between themselves over high data rate links. However, it was soon found that load-sharing introduces considerably more problems than it solves and the project was terminated. A similar experiment at Bell Labs has also demonstrated the infeasibility of load-sharing. Here, the project achieved reasonable success for several years before small discrepancies in the library routines at each location killed the effort.

Networks for message handling have been considerably more successful as demonstrated by the Westinghouse inventory control system and the airlines reservations networks, but the techniques utilized are special purpose and are not transferable into general intercomputer communication.

More recently, experiments have been carried out between Lincoln Laboratory and System Development Corporation to test the feasibility of more general computer-computer interaction. This experiment

demonstrated the relative ease of modifying time-sharing systems to permit network interactions and provided some statistics on the message lengths encountered. This experience has been added to through the introduction of the 338 display computer at ARPA tied to the Lincoln system. Although the requests in this communication link are totally in one direction, the form of communication utilized is identical with that expected in network activities and has extended the techniques to include graphic display interactions.

C. <u>Rationale</u>.

<u>Effect on ARPA Research Program.</u>

Seventeen computer research centers throughout the country are supported in whole or in part by Information Processing Techniques. The installation of an effective network tying these locations together should substantially reduce duplication and improve the transfer of scientific results, as well as develop the network techniques needed by the military. The research output of these projects is important to all three Services and it is expected that this output can be substantially increased for the same dollar cost if a portion of the funds are utilized for the network.

ARPA is in a unique position to carry on network research. Its computer projects involve most of the outstanding computer scientists in the country and these talents are necessary to the success of such a project. Further, since the projects are all centrally funded, the network can get off the ground without getting bogged down with problems of inter-project charging. Except for isolated experiments, such research could not currently be initiated outside ARPA.

D. <u>Plan</u>

1. <u>Preliminary Network Planning</u>.

In early 1967 preliminary plans for an interactive computer network were discussed with IPT contractors. Working groups were established to design standardized communications protocol and to specify their network requirements. A preliminary protocol was developed and discussed with interested parties during the summer of 1967. At that time planning for the network had progressed sufficiently for contractors to begin including tasks as they negotiated their Fiscal 68 contract renewals. Although extensive effort will not be required to modify system software at each network location, some modifications are necessary and have been planned for in this way.

In December 1967, a small contract (AO 1137) was initiated with Stanford Research Institute for the development of specifications for the necessary communications system. This effort has resulted in sufficiently detailed documentation to allow a request for bids on Interface Message Processors (IMPs). SRI will also provide continuing assistance to the initial participants in the network.

2. Network Information Center

In order for people to utilize the envisioned computer network effectively, it will be necessary to provide extremely good documentation on what programs and files are available throughout the net. This information should be available on-line to any individual in the network. It should be possible for him to add new program descriptions, edit previous descriptions, retrieve relevant information based on keyword searches and affix comments to program descriptions which he has used. To achieve this goal, Stanford Research Institute has been tasked with developing such a facility. This is an extension of the capability already achieved at SRI and is in progress in order that it may become available concurrently with the network.

3. Communication System

Multi-point, fast response, high capacity, reliable communications are required for an interactive computer network. The traffic between nodes is expected to consist mainly of short digital messages with a wide dispersal of destinations. Initially, message length will vary from one to one thousand characters with an expected average length of 20 characters. Since a cross country 50 kb communication line has a delay equivalent to 150 characters, messages must be continuously multiplexed into each line in order to maintain reasonable efficiency. Since the dispersion of destinations is large, messages with different origins and destinations must be concentrated into the same line. This can only be achieved with a store and forward system.

Message delay for on-line, interactive work should be well below one second (origin to destination). This cannot be achieved with voice grade communication lines in a store and forward system. However, with 50 kilobit communication lines, the required response speed can be attained. The additional capability obtained with 50 kb lines is also important, but is not the prime factor dictating the choice of these lines.

After considering the trade-offs associated with the communications subsystem, it was decided to design and build a store and forward net

using message processors at each research center interconnected with 50 kb communication lines. Such a distributed communication system will be revolutionary, providing vastly reduced transmission costs, fast response and high reliability. The effect of providing such an efficient communication capability to the computer researchers should be to inspire the development of creative and effective network techniques.

In order to develop an effective store and forward processor, a contractor will be selected competitively to take on the system responsibility for the communications subsystem. The following schedule outlines the expected events:

a. July 1968 -- Award IMP contract

b. March 1969 -- Demonstrate initial net operation with four nodes

c. April 1969 -- Approve design and extend contract to include installation of 19 IMPs

d. December 1969 -- Complete network operational

e. 1970 -- Add communication lines as necessary

f. 1971-1972 -- Arrange with a common carrier the transfer of the communications system

The funds required to implement this plan will not be known precisely until the RFP responses are received. However, it is estimated that the IMPs will cost \$50K each and that the contractor will require about \$40K per month. Thus the initial phase of the contract will require \$560K of FY 68 funds. The second phase will be funded from FY 69 funds and will require about \$1 Million. The communications will be procured separately from DECCO and for the initial net will cost about \$10K per month. As the network expands, the communications cost will rise to \$900K per year. This should provide communications at the remarkably low cost of 10^{-5} cents/bit.

4. Experimental Development of Network Techniques.

While the communications system is being installed, each research center will be making the necessary additions to its time-sharing system to accommodate network interactions. As soon as the communications system

becomes available, users will begin to experiment with remote languages and systems. Some activities have already been identified and interest will grow rapidly once reliable communications have been established. Researchers often transfer between the various projects and then must reprogram certain basic systems at the new center. With the network, these needs will be eliminated since the researcher may use his original program through the net.

5. Cost Estimates.

The interactive network costs fall into two categories: (1) The communications and Interface Message Processor costs, and (2) Costs of communication and resource sharing experiments. The majority of this second class of costs will be borne by each of the computer research contracts now extant. They will vary across a range of extremes bounded by, for example, a single researcher's small experimental program and a group of researchers concerned with studies of on-line documentation. The communications and Interface Message Processor costs are more easily identifiable. The table below shows the commitment requirements.

Year	Costs		
	Communication Line	IMP Contractor	Total
FY 68	0	563K	563K
FY 69	25K	1000K	1025K
FY 70	680K	200K	880K
FY 71	900K	100K	1000K

E. Transfer

The transfer of interactive computer network technology will occur in three forms: (1) Dissemination of techniques and experimental results through the open scientific and technical literature, (2) Through the common carriers or other commercial organizations concerned with data transfer and dissemination, and (3) Through the military command and control centers for which the National Military Command System Support Center in the Pentagon serves as the focal point.

1. Publication of Results.

Dissemination of new scientific and technical knowledge through conferences and the appropriate literature is a slow but

necessary process. It was through this means, and only this means, that computer time-sharing technology was transferred from ARPA research projects to private industry. Had representatives of computer hardware and software firms participated alongside these ARPA projects in the early months of the research, there is no doubt that the transfer would have occurred more quickly and efficiently. But these firms were unable (or unwilling) to make such a commitment until they learned of the projects' results through the various publicity media open to the projects. Thus it was not until four years after the beginning of time-sharing research and development that the computer industry announced their intent to market such systems. Two years following their announcements, the industry began to show some signs of understanding how to build such systems and are now beginning to deliver them in rudimentary form. Although this development cycle time may be of a shorter period than the traditional one, it is believed that it can be shortened further.

2. Communications Carriers.

The ARPA interactive computer network will serve as a test bed for both man-computer and computer-computer research experiments. The Interface Message Processors are specified to be programmable in order to facilitate network changes that may be dictated by desired experimental design changes. Because the network characteristics must be subject to change, it cannot provide the stable environment necessary to the support of a day-to-day, standardized, digital communication service. However, results from network experiments will provide design information for those commercial groups concerned with digital message services. Thus, when a stable network is desired, a common carrier can be requested to take over management of the IMPs and provide digital message service directly to the individual users on a tariff basis. This would permit ARPA to terminate its system responsibility. Furthermore, it will allow non-ARPA-supported users who are concerned with uses outside the domain of computer systems research to join the network, e.g., Office of Education Regional Laboratories, NSF-supported universities and various user groups supported by the NIH. All three of these agencies have expressed a desire to make use of the network at the time the focus shifts from computer systems research to stable user communications needs.

3. The Military Community.

Within the past year and a half Information Processing Techniques has created a close working relationship with the National

Military Command System Support Center. This is the primary computer organization serving the Joint Chiefs of Staff and operates under the management of the Defense Communications Agency. Its computer capability serves also as the fountainhead for military command and control centers around the world. IPT has paid for the development of a prototype time-sharing system to be tested in the NMCSSC over the next 12 months. As a result of these test findings, it is expected that the NMCSSC will move away from conventional batch processing and into the computer technology created by ARPA and its contractors. As it does so, other military centers affiliated with it will follow, e.g. CINCPAC, CINCEUR and MACV. Such a collection of affiliated military centers using computer technology represented by ADEPT-TDMS provides a natural recepient for an interactive computer network. Transfer will be facilitated by SDC's (creator of ADEPT-TDMS) participation in the ARPA experimental network. The ability of military command systems to be able to interactively call upon one another through a distributed net will be particularly important due to their joint responsibility and common data base interests.